CENTRE DE RECHERCHE
SUR LES TRANSPORTS
(C.R.T.)
CENTRE FOR RESEARCH
ON TRANSPORTATION

UNIVERSITÉ DE MONTRÉAL
ÉCOLE DES HAUTES ÉTUDES COMMERCIALES
ÉCOLE POLYTECHNIQUE

CENTRE DE RECHERCHE SUR LES TRANSPORTS
UNIVERSITÉ DE MONTRÉAL
C.P. 6128, SUCCURSALE A,
MONTRÉAL
CANADA H3C 3J7

TÉLÉPHONE: (514) 343-7575
TÉLEX: (21) 05-561325
EMAIL: crt@crt.umontreal.ca
TÉLÉCOPIEUR (FAX): (514) 343-7121

CRT-689

Design of Regular Intercity Driver Routes for the LTL Motor Carrier Industry

by
Teodor Gabriel Crainic[1,2]
Jacques Roy[1,2]

[1] Département des sciences administratives, Université du Québec à Montréal, C.P. 6192, Succursale A, Montréal, Canada H3C 4R2.

[2] Centre de recherche sur les transports, Université de Montréal, C.P. 6128, Succursale A, Montréal, Canada H3C 3J7.

Centre de recherche sur les transports — Publication #689

May 1990, Revised November 1991

92-02305/

Dépôt légal – 4ème trimestre 1991

ABSTRACT

The load plan of a LTL motor carrier determines the services it operates and the frequencies it offers. To satisfy the stable part of the demand for transportation, regular driver routes have to be built. These routes must cover the planned services and abide by the laws, regulations and conventions that regulate the drivers' workload, while guaranteeing the timely delivery of goods at the destination terminal. The objective is to minimize costs and to ensure an efficient fleet utilization. We describe the problem, examine its characteristics, and propose a modeling and algorithmic approach based on set covering ideas. Results of an actual application are also presented.

KEY WORDS: LTL Motor-Carrier Transportation,
Routing and Crew Scheduling,
Regular Driver Routes,
Set Covering.

RESUME

Le plan de transport d'une entreprise de camionnage de charges partielles détermine les services et fréquences à offrir. Pour assurer le service de base et afin de répondre à la partie stable de la demande, il faut établir un ensemble de routes–horaires régulières qui couvrent les services et qui respectent les lois, règlements et conventions régissant le travail des chauffeurs, tout en assurant la livraison à temps des marchandises au terminus de destination. L'objectif est de minimiser les coûts d'opération et d'assurer une utilisation efficace de l'équipement. Nous décrivons le problème, en examinons les caractéristiques et présentons une approche de modélisation et de résolution basée sur des techniques de recouvrement d'ensembles. Les résultats des expérimentations effectuées à partir d'une application industrielle sont également présentés.

MOTS CLES: Transport routier de marchandises en charges partielles,
Etablissement de routes et horaires,
Routes–horaires régulières,
Modèles et méthodes de recouvrement.

INTRODUCTION

Less-than-truckload (LTL) motor carriers often design *regular driver routes* in order to satisfy that part of intercity freight transportation demand which is known to have a high degree of certainty. This relatively stable part of the demand for LTL services may account for approximately 60% to 80% of total demand, depending on seasonal variations. In this way, carriers are able to provide their customers with a reliable service, while offering their drivers a choice among a set of regular routes. In many cases, according to their seniority in the company, drivers may bid for the routes they will be driving for the next six months or so. The use of such regular routes also enables motor carriers to improve their control over the utilization of equipment and facilitates dispatching operations.

In order to properly design regular routes, it is usually necessary to prepare ahead of time a transportation (load) plan that will take into account the ever present trade-offs between customer service performance and operating costs at a tactical planning level. This *transportation plan* generally determines:

- the type of transportation services to be operated between each pair of terminals and their frequencies;
- the itineraries (sequence of services and terminals) that are to be used to move freight from its origin to its destination;
- the workload, in terms of tonnage handled and number of trucks loaded and unloaded, at each terminal in the network;
- the necessary empty vehicle movements required to establish network balance.

Regular driver routes are then designed in such a manner as to cover those services derived from the tactical transportation plan, while respecting both government and labour contract rules. In this case, the objective is the on-time delivery of freight to its destination, while minimizing operating costs.

The major difficulty nevertheless lies with the efficiency of the regular routes. Indeed, costs related to different routes will vary according to:

- the extent to which productive driver hours are actually used, the total amount of permissible production time being restricted by government regulation and labour contract rules;

- the amount of time required to hook up and change trailers;
- the time spent for general maintenance, fuel replenishment, various checks, etc.;
- the use of overtime;
- off-base allowances, such as meals, lodging, etc.

When solving such a difficult problem manually, managers will typically consider only a few alternatives in their analysis. They will most often be satisfied with a few improvements to the solution that was used during the previous planning period. The whole exercise will depend on the manager's experience and knowledge of the problem. However, such an approach will not always provide motor carriers with effective and efficient regular driver routes. This is particularly true in the present context of rapid market changes and more restrictive safety regulations that affect drivers' time on the road and rest periods. Indeed, such regulation has just recently been enforced in the Province of Québec, thereby severely restricting the permissible length of direct links when designing driver routes. Yet, to our knowledge, no comprehensive research effort has been directed toward this issue of designing efficient regular driver routes for the LTL motor carrier industry.

The objective of this paper is to describe this problem and identify its main characteristics. We present a general solution approach that ensures the efficient and "optimal" generation and selection of regular driver routes. The emphasis is on the formulation of a model that adequately reflects the operational and decision-making structure of the LTL motor carrier industry and that may be solved efficiently. We will also illustrate the substantial performance improvements our approach may yield for motor carrier operations.

The paper is organized as follows. Section 1 presents the problem, its general planning context, and the main components of our solution approach, while in Section 2 we specify the notation and some necessary definitions. Section 3 presents the modeling of the route cost and productivity measures. Section 4 deals with the problem of generating segments. Section 5 is dedicated to the model and algorithm used to generate and select the best set of regular driver routes. Section 6 presents results obtained with real data, and we conclude with a discussion of further research and development issues in this area.

2. GENERAL PROBLEM AND SOLUTION APPROACH

Figure 1 shows a schematic representation of the planning process which leads to the design of regular intercity driver routes in the LTL motor carrier industry. This planning process is implemented, in various forms, by most motor carriers using regular driver routes.

Figure 1. General Planning Approach for the Design of Regular Intercity Driver Routes

First, forecasts must be made for that part of the demand for transportation services which is considered to be stable over a given planning horizon (or season). Indeed, regular driver routes are only useful if a given level of demand can be determined with a high degree of certainty. For major Canadian LTL carriers, this stable part of demand may represent from 60% to 80% of total demand, depending upon the season.

Forecasts must be made for each market served by the carrier. A *market* is determined by the origin and destination of the traffic, normally a pair of cities (terminals). Since we are only concerned with freight moving between existing terminals within a given motor carrier's network, we assume that demand is specified by the amount of freight to be moved from an origin to a destination terminal, for a given period of time.

These market forecasts are then used as inputs to the *tactical planning process* [7], which produces a *transportation* (or *load*) *plan* for the stable part of demand for LTL services. Such a plan can be established, for example, by using the NETPLAN software [19], developed from a tactical planning model proposed by Roy[18] and based on a methodology put forward by Crainic[3,4] and Crainic and Rousseau[5], or the interactive optimization system [16] based on ideas proposed by Powell and Sheffi[15] and Powell[14]. Although such advanced methods are being used by a number of major U.S. motor carriers [16], manual methods are still, however, widely used.

Services and *itineraries* are important elements of the tactical planning process. A transportation *service* is determined by a transportation mode (e.g. trailer type) that can be used between two terminals within the network. A service is also defined by the *path* it follows through the physical network from its origin to its destination. Two services are said to be *symmetric* if they operate between the same two terminals and on the same path but in opposite directions. Although the path characteristic is often neglected at the tactical planning level, it plays a very important role when it comes to designing intercity driver routes, as we will see later on. *Frequency* is another important aspect of a service; it determines the number of departures, i.e. the *level of service*, associated with a particular transportation service for a given period. The set of all transportation services that a carrier makes available to its customers defines its *service network*.

For a given month, an *itinerary* defines how freight will flow from origin to destination. It determines the sequence of services used to move the freight, and the type of operation performed on freight at each intermediate terminal. Although itineraries play an important role in defining the transportation plan, as do services, they do not directly affect the design of regular driver routes; services, however, do. Indeed, the type and frequencies of the transportation services used to satisfy the stable part of demand will be determined through the tactical planning process, and will then be used as input data to the regular driver route design

process, as illustrated in Figures 1 and 2.

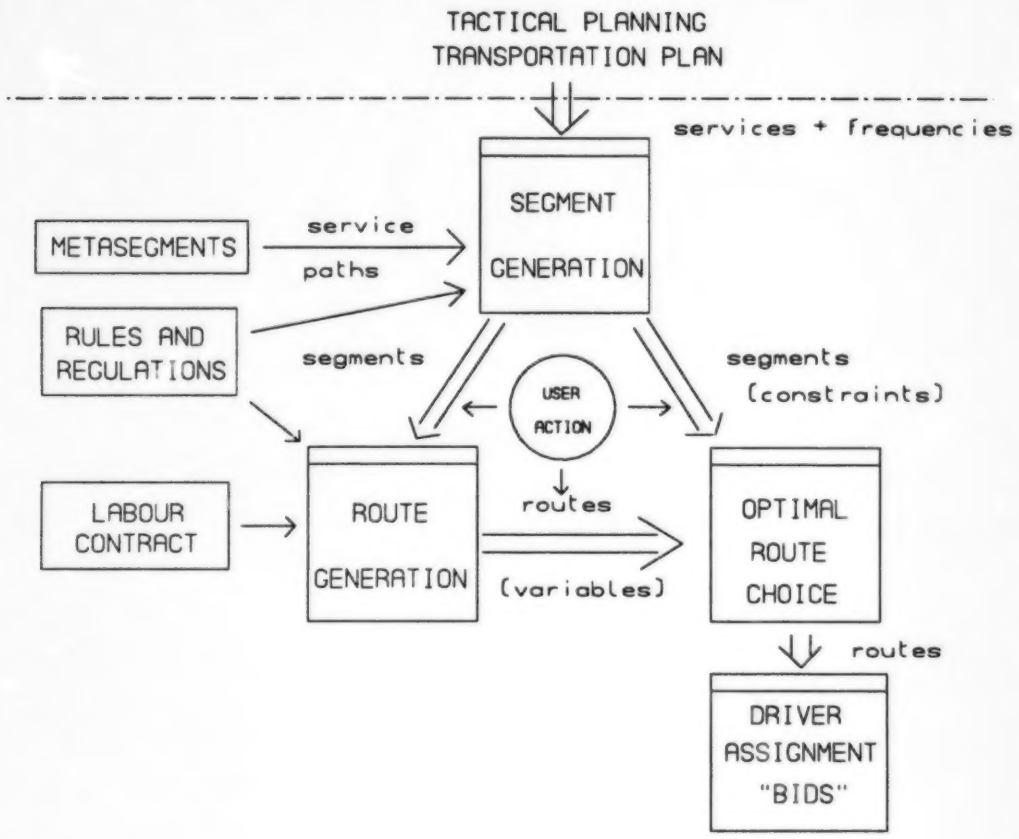

Figure 2. General Design of the Solution Approach

We did not find any published material dealing with the problem of the design of regular driver routes. An examination of the characteristics of the problem, especially in comparison with crew scheduling problems in other transportation modes [17], allows us, however, to identify certain similarities with the airline crew scheduling problem [6]. We have made use of these similarities in developing our solution method.

We propose a solution approach, illustrated in Figure 2, which starts with the transportation services and frequencies specified by the tactical load plan, and which takes into account both government and company rules and regulations. The problem is dealt with in three stages:

(i) *Segment Generation*. The services specified in the transportation plan must first be broken down into a set of *segments*. Each segment represents a direct driving link, and each service path yields one or several segments.

(ii) *Route Generation*. In order to satisfy the stable part of demand, the segments generated in the first stage of the procedure must be covered by regular driver routes. Such a set of efficient legal routes is generated at this stage.

(iii) *Route Selection*. We propose a model and solution approach that will select an optimal set of regular routes which satisfy the demand expressed by the segments generated above.

Routes from this optimal set are then assigned to drivers on a seniority basis as part of a bidding process. The following sections deal with each one of these methodological components.

2. DEFINITIONS AND NOTATIONS

Let us define the *planning period*, or simply the *period*, as the span of time, usually one week, for which regular driver routes have to be specified. These routes will then be repeated for all subsequent periods. Starting times for all components of the problem are computed with respect to the beginning of the planning period, which thus represents the "zero hour".

A period is usually broken down into several days. A *day* corresponds to a cycle of work and free time intervals (including rest). In intercity trucking, however, a working day does not normally coincide with a typical calendar day: a driver's workday starts in the evening and finishes the following morning. We should therefore refer to a driver's working period as a "work night". However, for the sake of clarity, we shall use the more conventional concepts of workday and rest periods.

Our approach to the problem is based on three major concepts: the *service path*, the *segment* and the *driver route*. These elements may be defined according to various attributes. Some of these are descriptive: location, length, etc. Others are time related and indicate when a particular movement occurs. We call this moment the *starting time*. It is important to note that an element has only one physical description, but that it may have several time occurrences. For example, one could cover the same path several times during a given planning period. To distinguish between descriptive and time-related attributes, we use the prefix *meta* whenever the temporal attribute of an element is irrelevant.

Motor carrier operations are best described in terms of various *networks*. These networks are all defined by a common set of nodes; the link sets may vary.

A *node* is a geographical location, where an operation may be performed. It is either a *terminal* or a *transfer point*. A trailer exchange may normally be performed at any node, but no service may begin or end at a transfer point. A *base* is a terminal where drivers are stationed. Figure 3 provides an illustration of these definitions. Nodes 1, 2 and 3 represent terminals which are also bases, while node 4 is only a transfer point. We use B as the set of all bases, T as the set of all terminals and N as the set of all nodes, with $B \subseteq T \subseteq N$.

Define d_j^{op} as the *minimum delay* required to perform the specified operation *op* at node j, $j \in N$. At transfer points, this operation may only consist of an exchange of trailers

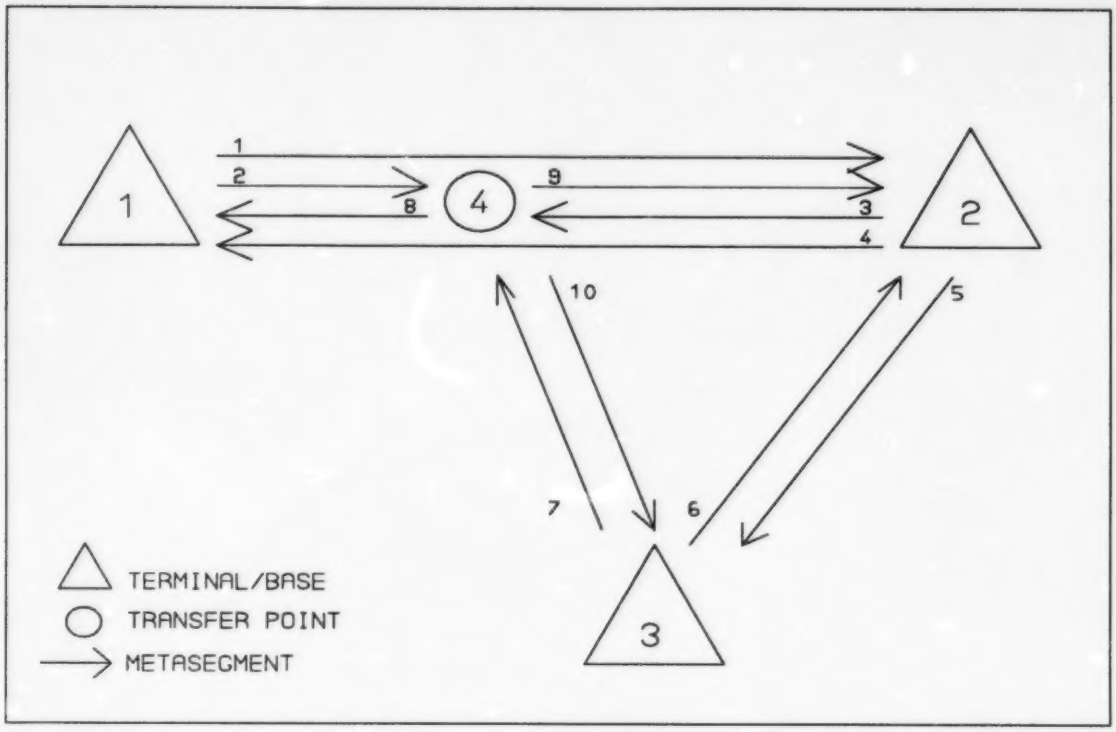

Figure 3. Network of Nodes and Metasegments

between two symmetrical services. At terminals, however, the driver may either perform an exchange of trailers, or drop off a trailer of a given service and pick up a trailer for a new destination. Note that, in practice, the same minimum delay tends to be specified for each type of operation, pick up, drop off and transfer, irrespective of the particular node.

A *metasegment* represents a direct link between two nodes in the network which may be accomplished, according to company rules and regulations, without any intermediate operations. Let Σ be the set of all metasegments. For a given metasegment s, $s \in \Sigma$, we define:

- *Origin* O_s, $O_s \in N$;
- *Destination* D_s, $D_s \in N$;
- *Average driving time* δ_s;
- *Frequency* f_s: the number of times that the metasegment is used by the services specified in the current transportation plan during the planning period.

The metasegments define a network of permissible movements $G_\Sigma = (N, \Sigma)$. For example, in Figure 3, a total of ten metasegments determine the movements allowed between the nodes of the network. In this case, the absence of metasegments between nodes 1 and 3 does not necessarily imply an absence of roads between these two points, but may simply reflect a particular company policy.

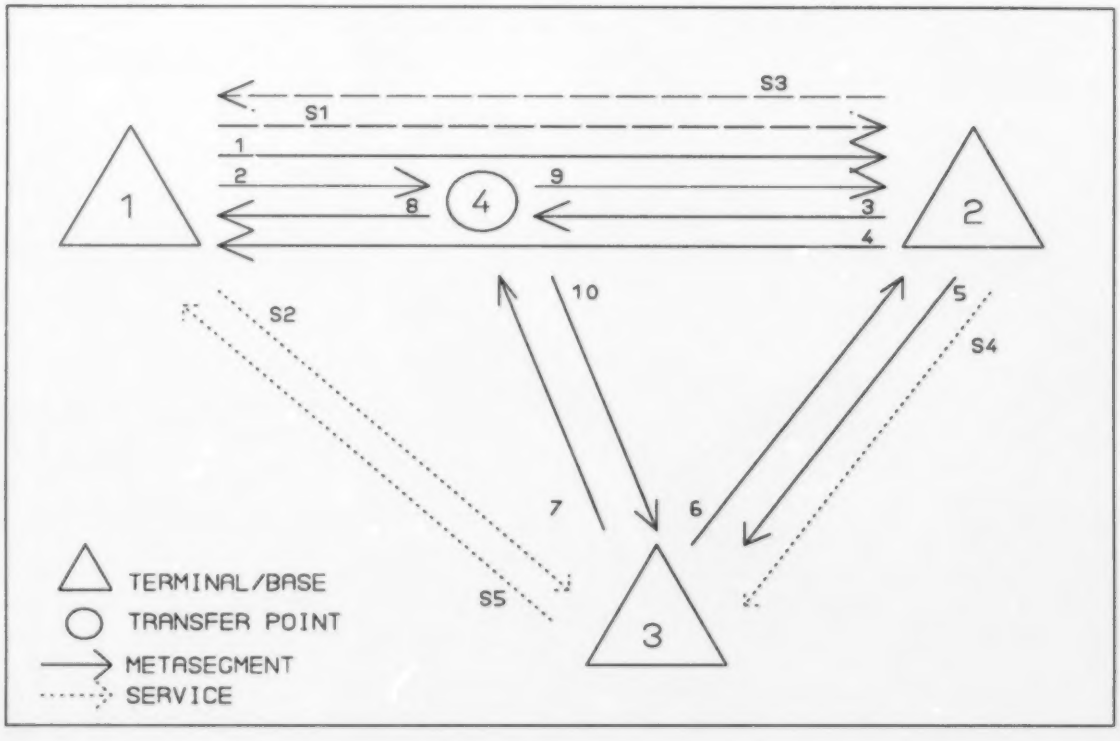

Figure 4. Network of Services and Metasegments

The metasegment network also determines the various service paths which may be used to implement the transportation plan. Thus, let us suppose that five services, illustrated in Figure 4, are to be operated on the network and covered by the metasegments previously defined. Several service paths may then be used, as illustrated in Figure 5. In particular, we observe that a single service S1, for example, may use several paths, such as S1.1 or S1.2 and S1.9. On the other hand, we also note that a given metasegment may be part of several service paths, such as metasegment 2 for services S1 and S2, 3 for S3 and S4, etc.

A *segment* $s, s \in S$, is a metasegment to which a temporal dimension is added by the specification of a *starting time* $t(s)$. Hence, several segments, such as $(O_s, D_s, t_1(s))$,

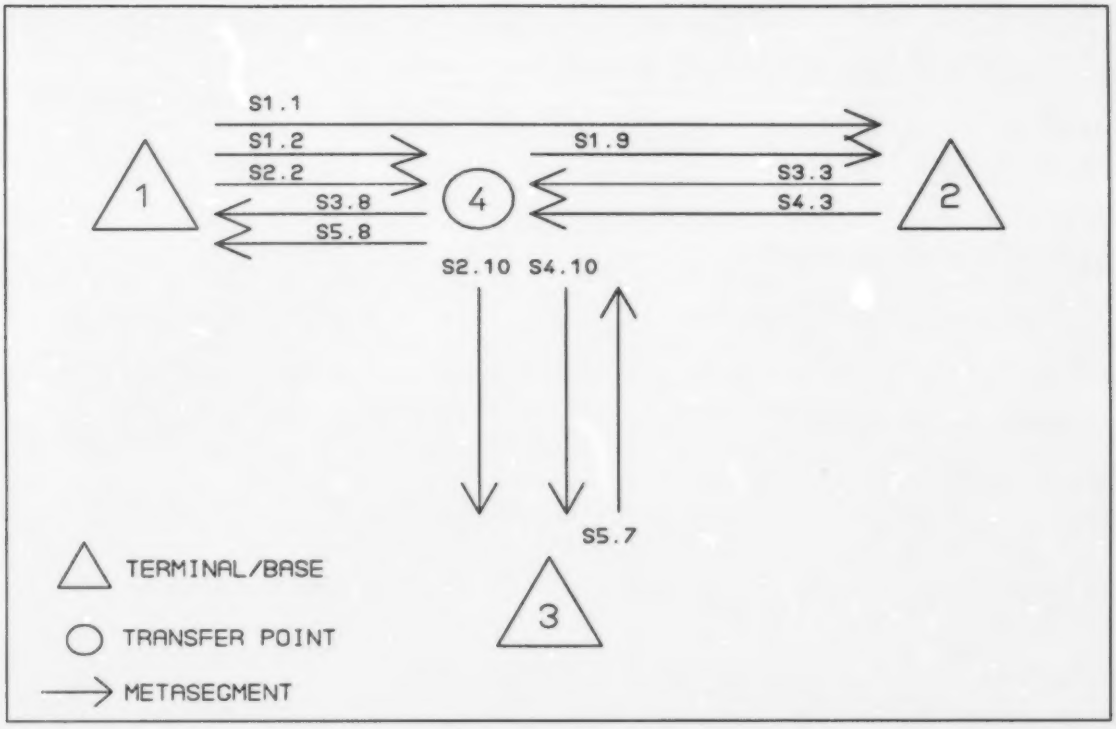

Figure 5. Network of Service Paths

$(O_s, D_s, t_2(s))$, ..., $(O_s, D_s, t_n(s))$, could correspond to a given metasegment (O_s, D_s) each segment individualized exclusively by its own starting time. S represents the set of all segments.

Segments and metasegments represent the basic units which make up the service paths and driver routes. They therefore play a major role in our model (constraints are derived from these elements) and in the procedures aimed at service analysis and generation of driver routes.

A *metaroute* is a legal sequence of metasegments, which starts and ends at the same base. The metaroute is said to be *normal* if it does not include any rest period outside its base, and *extended*, otherwise. A metaroute whose composing metasegments appear only once is called *simple*; otherwise, it is *cyclic*. Table 1 displays three metaroutes built by using the metasegments illustrated in Figures 3 and 4. Note that the metasegment (2,2) represents a rest period at node 2.

DESCRIPTION	TYPE
(1,4) (4,3) (3,4) (4,1)	NORMAL – SIMPLE
(1,2) (2,2) (2,4) (4,1)	EXTENDED – SIMPLE
(1,4) (4,1) (1,4) (4,1)	NORMAL – CYCLIC

Table 1. Examples of Metaroutes

Several attributes define a metaroute r, $r \in \Omega$, where Ω is the set of all metaroutes. Those required by our formulation are:

- its *base* b_r, $b_r \in B \subseteq N$, which is both the origin and the destination of the metaroute;
- its *composition* κ_r, which is made up of a sequence of metasegments (hence, nodes) and, eventually, off-base rest periods;
- the *driving time* δ_r, which is computed as the sum of the average driving times for all the metasegments that make up the metaroute: $\delta_r = \sum_{s \in \kappa_r} \delta_s$;
- the *duty time* w_r, which represents the amount of time the driver is available for work, either driving a vehicle or waiting for the completion of an operation at a node: $w_r = \delta_r + \sum_{j \in \kappa_r} d_j^{op}$; note that the minimum delays for the first pick up and the last drop off operations (both occur at the base) are increased to account for the preparation time at the beginning of the route, and for the time required to report in at its end;
- the *total trip time* τ_r, which adds up the length of the rest periods to the duty time; of course, for normal routes $\tau_r = w_r$;
- its *cost* c_r, which is defined in Section 3.

A *driver route* (or simply, a *route*) is a legal sequence of segments starting and ending at the same base. It is therefore a metaroute to which a time dimension has been added through the specification of the starting time of each constituent metasegment. Hence, a route has a specific *starting time*, $t(r)$, which corresponds to the starting time of the first segment in its composition. With the appropriate adjustments, all other definitions and notations relative to metaroutes also apply to routes. In particular, the composition of a cyclic route displays at least two occurrences of the same metasegment, at different times. On the other hand, the duty time of an actual driver route is longer than the duty time of the corresponding

metaroute, since it has to account for all the additional delays (on top of the minimum delays d_j^{op} already discussed) due to waiting while the trailer is being prepared (loaded), or waiting for a specified departure time. These additional delays are indicated by the starting time of the "next" segment.

These routes will then be assigned to drivers, according to their seniority in the company. Hence, they correspond to the decision variables of our model.

3. COST MEASURES FOR DRIVER ROUTES

"Cost" is the most widely used productivity measure for evaluating driver routes, and drivers' wages form its major, often its sole, component. Yet, this measure is not sufficiently precise, nor sufficiently discriminating, for an optimization-based procedure which selects the best set of regular driver routes. Hence, we present in this section a productivity measure which, while still based on the wages drivers receive, is better suited to the requirements of the route selection procedure.

Drivers' wages are normally regulated by labour contracts which guarantee for each day (duty period), that a minimum number of duty hours shall be paid, irrespective of the actual amount of time worked. The wage formulae also specify the maximum spreads for normal duty and overtime, as well as the hourly rates for the two cases. A realistic example of such a formula might indicate that the equivalent of 9 worked hours are guaranteed for each worked day, that the maximum normal duty period is of 12 hours, including breaks, and that the driver may work a maximum of 14 hours including overtime. Various other factors, such as distance rates (which vary according to the route), route dependent overcharges, and auxiliary expenses (e.g. monetary compensations for rest periods taken off-base), further complicate the accounting formula for driver costs. Blanchette[2] provides a detailed presentation of the procedure used to compute drivers' wages for a Canadian LTL motor carrier. The similarity between this procedure and the methods used by North American airlines is remarkable.

A realistic representation of the time dependent part of the *real* cost γ_r of a normal regular driver route is:

$$\gamma_r = \begin{cases} K^m & \text{if } 0 \le w_r \le \omega^m \\ K^m + c_r^n\,(w_r - \omega^m) & \text{if } \omega^m \le w_r \le \omega^n \\ K^m + c_r^n\,(\omega^n - \omega^m) + c_r^o\,(w_r - \omega^n) & \text{if } \omega^n \le w_r \le \omega^o \end{cases}$$

where

ω^m : Minimum guaranteed duty time, per day;

ω^n : Regular (normal) duty time spread for a day;

ω^o : Maximum duty time in a day, including overtime;

K^m : Minimum guaranteed wage;

c_r^n : Hourly rate for normal duty time, for route r;

c_r^o : Hourly rate for overtime, for route r;

and the route is, of course, legal: $0 \leq w_r \leq \omega^o$. The cost of an extended route is computed as the sum of the costs of each of its duty periods, obtained by using this formula, plus the costs of the auxiliary expenses (e.g., the cost of off-base rest periods).

This formula **is not**, however, appropriate for computing cost coefficients for the objective function of an optimization-based route selection procedure, since it does not adequately discriminate among routes with similar cost but different productivities. Indeed, on the one hand, any route with a total trip time less than the guaranteed minimum receives the same cost, K^m, irrespective of the actual work done. On the other hand, this type of formula does not identify the operations that make up the cost of the route. Hence, a route with long inter-service delays may receive the same cost as a route where most of the duty time is dedicated to driving. Yet, if one agrees that *driving* the truck represents the *productive* part of a driver's work assignment, one would prefer to select routes where as few as possible of the paid hours correspond to non-driving activities.

To correct these shortcomings, we define the cost of a driver route as the *time paid but not dedicated to the driving of the vehicle*, and use this measure for the route generation and selection methodology described in this paper (see the Appendix for a brief discussion of alternate formulae). The *paid time* for a normal route, $\pi_r^{\$}$, may be defined as follows:

$$\pi_r^{\$} = \begin{cases} \omega^m & \text{if } 0 \leq w_r \leq \omega^m \\ w_r & \text{if } \omega^m \leq w_r \leq \omega^n \\ \omega^n + \dfrac{c_r^o}{c_r^n}(w_r - \omega^n) & \text{if } \omega^n \leq w_r \leq \omega^o \ . \end{cases}$$

The cost of a normal route is then $c_r = \pi_r^{\$} - \delta_r$. The cost of an n-day extended route is computed by a similar formula. First, the paid time of each one of the n duty periods $\pi_i^{\$}$, $i = 1,...,n$, is computed by using the formula defined above. The cost of the extended route is then given by $c_r = \sum_{i=1,...,n} \pi_i^{\$} + nK^{ob} - \delta_r$, where K^{ob} represents the cost of an off-base rest period.

The use of this formula implies that all components and activities (e.g. overtime periods, off-base rest periods, etc.) that make up a driver route are first evaluated in terms of normal

driving time. The conversion factor between normal and overtime working hours may vary. The ratio between the two corresponding cost rates seems appropriate since it reflects the monetary relationship between the two activities. One may also use increasing conversion factors for each hour (or portion thereof) of overtime duty. Then, the paid time is compared with the actual driving time, and the difference is assigned as the cost of the route. Thus, short routes are marked as "more expensive", as are those with extensive overtime periods, or long waiting periods at terminals.

It is noteworthy that this formula displays characteristics which are similar to those of the formulae used to compute the crew pairing cost in most North American airlines (see [6] or [12], for example). As in these formulae, the regular driver route cost formula we propose is intrinsically nonlinear: the cost variation due to the addition to, or the deletion of, a segment from a route, depends not only on the specific segment, but also on the position of the segment in the route path and on the route's total trip and duty times. Consequently, one cannot evaluate the cost of a segment **prior** to the generation of the route. The segment and route generation procedures are discussed in the next sections.

4. SEGMENT GENERATION

Segments represent the basic data for the whole process. Driver routes, metaroutes and service paths are all constructed by using segments, whether they are time related or not. Moreover, segments are also used to generate the constraints of the mathematical formulation proposed in the next section to optimize route selection.

To generate segments means to determine which metasegments, from among those the company has specified as "feasible", are to be used during the contemplated planning period, and to specify their starting times, while insuring on-time delivery of freight at destination, minimizing total times and reducing the number of off-base rest periods.

Conceptually, segments are obtained from the service paths and frequencies specified in the (tactical) transportation plan. The path of each service departure is broken down into metasegments, each of which generates a segment once its starting time is computed. This computation is carried out by adding the driving time associated with each preceding metasegment, plus the delays at nodes, to the time the service starts from its base.

This approach implicitly assumes however that the starting time and exact path of each service departure are known ahead of time. However, this information is either unavailable, or unreliable. In fact, this data is not required for the preparation of the transportation plan at the tactical planning level. Furthermore, when it comes to implementing this plan into day-to-day operations, dispatchers may significantly modify the service departure times in order to respond to variations in the actual demand for transportation services. Indeed, the exact coordination of daily departures is performed in real time, the driver being advised only of his precise departure time within a relatively wide time interval (in the order of three hours, according to Blanchette[2]).

We need, therefore, to estimate an approximate starting time for each of the planned service departures that correspond to the stable part of demand. Our approach deals with each service separately and consists of two steps. First, we determine the number of daily departures required according to the service frequency specified in the transportation plan for the given planning period. Next, these departures are distributed, for each day of the period, over a time interval which, on one hand, allows sufficient freight to arrive and be handled prior to the first departure to ensure that the trailer is fully loaded and, on the other hand, schedules

the last departure early enough to enforce on-time delivery at destination.

In the first step, frequencies are distributed either evenly over the number of working days of the planning period, or according to information inferred from the historic pattern of service departures (which is significant since we are working with the stable part of demand). For the second step, experience has shown that departures are not necessarily evenly spread over the interval; very early and late departures can be observed, with most departures clustered during "peak" hours. When known, we use this pattern to estimate departure times. We also tested a uniform distribution of departures (see Section 6), and obtained good results, once symmetrical services and their corresponding metasegments were given special treatment: the time interval between the earliest and the latest departures is split in two in order to account for possible connections (round trips) between segments. At the same time, we increase the flexibility of the route generator, by introducing *overlapping periods* (Section VI), to better capture the scheduling flexibility inherent in the motor carrier industry and to take into account the necessary coordination required to transfer trailers from one segment to another.

When service paths have to be determined, we propose a method which makes use of shortest path techniques, and which is based on the observation that LTL companies tend to avoid off-base rest periods as long as deliveries may be made on time. This mode of operations is illustrated in Figure 4 where, for example, one notes that service S1 may be operated on two different paths: the first path uses metasegment 1 only, while the second travels on metasegments 2 and 9. The first alternative is faster, but it forces the driver to spend a rest period off base. The second path is more convenient if we operate as follows: a driver leaves node 1 with a trailer for node 2, while another driver leaves node 2 with a trailer for node 1. At the assigned transfer point, here node 4, trailers are exchanged, and each driver heads for his own base, his tractor pulling the trailer from the other service. With this approach, off-base rest periods are eliminated at the expense of a slight increase in the total transit time of each service and in the total duty time of each driver. In actual operations, the second alternative is used for early departures, thus permitting on time deliveries, while the first path has to be used for late departures if target delivery times must be met.

The method works on the metasegment network G_Σ, and penalizes long metasegments which would imply off-base rest periods, as long as the delivery deadlines can be met. Each

arc $s \in \Sigma$ has a *normal* cost $\delta_s + d^{op}_{D_s}$, and a *penalized* cost $\delta_s + d^{op}_{D_s} + K^{ob}$ if it is "too" long (i.e., if $\delta_s \geq \omega^o/2$); the penalized cost equals the normal cost for all other metasegments. Then, for any given service, on a given day, the algorithm is:

Initialization.

- Consider a service (O, D), with a daily frequency f,
  $t(i)$, $i = 1, ..., f$, the planned starting times, and
  L_{max} the latest delivery time at the destination D.
- $i = 1$.

Step 1.

- Find the shortest path from O to D in G_Σ with the **penalized** arc costs;
- Let L be the duration of this path.

Step 2.

- While $i \leq f$ and $t(i) + L \leq L_{max}$:
  - Consider this shortest path as the service path;
  - Decompose the path and compute the segments' starting times;
  - $i = i + 1$.

Step 3.

- If $i > f$ STOP; Otherwise
  - Find the shortest path from O to D in G_Σ with the **normal** arc costs;
  - While $i \leq f$
    - Consider this shortest path as the service path;
    - Decompose the path and compute the segments' starting times;
    - $i = i + 1$.
  - STOP

This approach allows efficient access to a wider array of solutions than simply resorting to paths that were used historically by the company.

5. IDENTIFYING THE OPTIMAL ROUTE SET

In a combinatorial sense, a driver route is any legal arrangement of segments, waiting periods at terminals and, eventually, off-base rest periods. Consequently, the total number of legal driver routes for a given case could be very high, especially when extended routes are contemplated. To generate all these routes would be both costly and algorithmically inefficient: the final set would be far too large for any efficient route selection procedure. Our approach aims at generating only a limited number of "good" regular driver routes. This set is "good" in the sense that it contains the optimal solution to the optimization problem described hereafter.

To achieve this objective, we propose a *set covering* approach in order to model the problem of selecting the best set of regular driver routes. Set covering is a classical modeling framework ([20], for example), quite familiar in the crew scheduling area [17], for which numerous solution methods have been proposed (e.g. [11] and references herein). To solve our model, we use a column generation based methodology, a technique which has been shown to offer extremely efficient solution procedures for complex, usually very large scale, crew scheduling problems ([1], [5], [6], [9], [10], [13], etc.).

We first state the model and the general algorithm. We then discuss the issue of generating routes and the role of the column generation technique in this context.

5.1 The Model and Algorithm

Consider R, the set of all regular driver routes. We may rewrite the composition (κ_r) of a route r, $r \in R$, as a vector $a_r = (a_{sr}) \geq 0$, such that $a_{sr} = 1$ if segment s belongs to route r (and equals 0 otherwise). Vector a_r is of dimension $n_S = |S|$, the cardinality of set S.

Let $x_r = 1$ if route r, $r \in R$, is chosen (and 0 otherwise). The set covering model that we use to select the optimal set of regular driver routes may then be written as

$$(P) \qquad \textit{Minimize} \sum_{r \in R} c_r x_r$$

$$\textit{Subject to} \sum_{r \in R} a_{sr} x_r \geq 1 \qquad \forall\, s \in S$$

$$x_r \in \{0, 1\} \qquad \forall\, r \in R\ .$$

We propose a set covering formulation, rather than a set partitioning approach (which ensures that each segment is selected only once in the optimal set), since, in practice, one may observe movements of tractors without trailers (rather rare), as well as regular LTL drivers hauling "irregular" LTL loads or full loads. Movements belonging to these last two categories occur quite frequently due to imbalances between supply and demand at terminals (especially the main ones) of the network. Hence, a set partitioning approach may be infeasible for most occurrences of this problem.

In order to determine the optimal solution to the above formulation, we first solve its linear relaxation (PL), where non-negativity constraints, $0 \leq x_r \leq 1 \ \forall \ r \in R$, replace the integrality conditions of (P). To do this, we make extensive use of the very interesting structure displayed by the constraint matrix of the model which is shown in Figure 6 where the columns correspond to routes and the rows to segments. Note that routes have been grouped in such a way that the normal routes for each day are shown together, and that segments have been sorted accordingly.

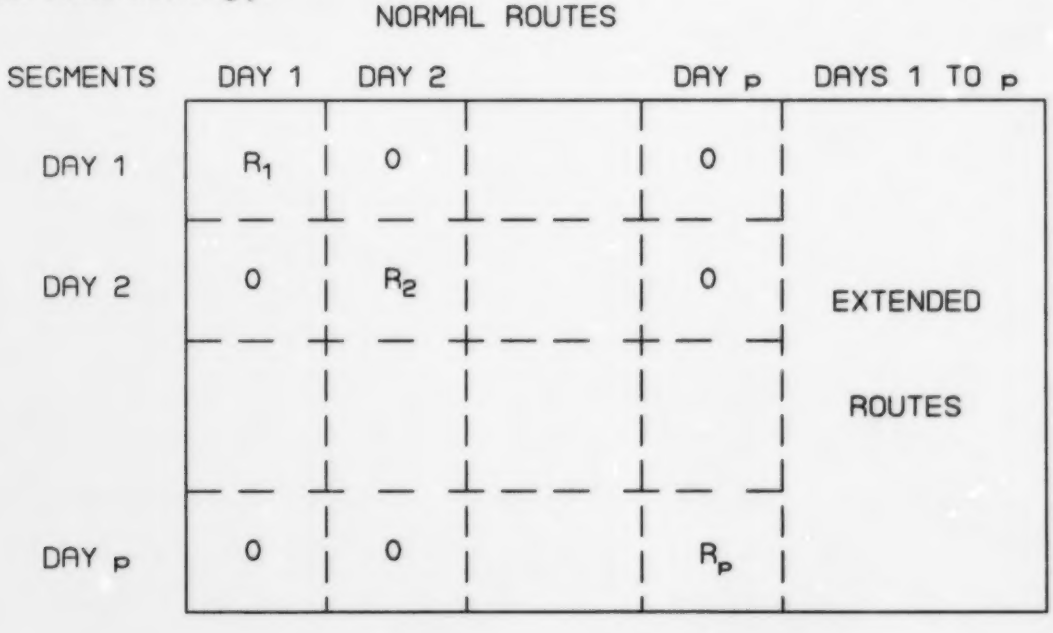

Figure 6. Structure of the Constraint Matrix

Figure 6 illustrates the fact that the set of all segments may be divided into p mutually exclusive subsets, where p represents the number of periods or "days" that normally make up the planning period (usually 5):

$$S_j = \{\ s \in S \mid \text{segment starts at period } j,\ 1 \le j \le p\ \}$$

$$S = S_1 \cup S_2 \cup \ldots \cup S_p \text{ and } S_j \cap S_i = \emptyset\ \ \forall i,j\ \ 1 \le i,j \le p.$$

The set of regular driver routes, R, may similarly be split into $p+1$ subsets R_1, R_2, ..., R_p, R_e, where

$$R_j = \{\ r \in R \mid \forall s \in \kappa_r, s \in S_j; \text{ i.e., normal route }\}\ \ 1 \le j \le p$$

$$R_e = \{\ r \in R \mid r \text{ extended route }\}.$$

Thus, the constraint matrix naturally decomposes into two components, one of which constitutes a block diagonal matrix, each block corresponding to the constraint matrix of a one-day-segments-normal-routes-only set covering problem. Therefore, we use the following algorithm to solve the set covering problem:

Step 1. For $j = 1, \ldots, p$

Solve the linear relaxation (PL) of the set covering subproblem which corresponds to period j (constraint matrix formed by segments in S_j and normal routes in R_j);

Let R_j^* be the set of routes which form the optimal solution to this problem and let Π_j^* be the corresponding simplex multipliers;

$R^* = \bigcup_j R_j^*$ and $\Pi^* = (\ \Pi_1^*\ \ \Pi_2^*\ \ \cdots\ \ \Pi_p^*\ )^T$.

Step 2. STOP if R^* covers all segments in S; Otherwise

Let $\bar{S}$ be the set of segments not covered by routes in R^*.

Generate, by using the Π^* multipliers, extended routes, $\bar{R}$, to cover segments in $\bar{S}$.

Reoptimize; let R^* represent the set of routes in the optimal solution to the linear relaxation (PL).

Step 3. Extract, from R^*, R^I, the integer optimal solution to problem (P).

Reduction rules could be added to this procedure in order to decrease the size of the covering problem and thus, hopefully, increase the efficiency of the solution method. Crainic and Roy[8] present several such rules as part of an extensive analysis of the computational performance of the set covering approach in the context of scheduling LTL drivers. Their study shows, however, that reduction rules to not significantly contribute to the efficiency of the

algorithm. Hence, the method described herein does not include them.

Note that, given a set of simplex multipliers Π^*, the marginal cost of any route can be computed by using the "classical" linear programming formula $\bar{c}_r = c_r - \Pi^{*T} a_r$ which, in the present context, is equivalent to $\bar{c}_r = c_r - \sum_{s \in K_r} \pi_s^*$, where π_s^* is the simplex multiplier associated with segment (constraint) s, $s \in S$. It is therefore extremely easy to compute the reduced cost of a route, given its composition and the current optimal solution of the relaxed problem. Consequently, the relaxed subproblem *(PL)* can be solved very efficiently with a combination of post-optimization and column generation techniques. We further discuss this procedure in the following section.

The decomposition approach that we propose is motivated by the observation that potential extended routes are by far the most numerous, and that, at least in the Canadian context, extended routes imply extra costs (notably for off-base rest periods) which make most of them rather uninteresting. Several factors support this second observation, including the relatively high density of the network of LTL terminals, the numerous, easy and inexpensive opportunities for implementing driver relief points (this is a major difference when compared with the airline crew scheduling problem), and the recent strengthening (in Canada) of the safety legislation concerning duty and trip times.

More significant, however, is the fact that the drivers' wage formula is basically a "per day" accounting operation, which simply adds up the duty times of the periods making up a route and the corresponding off-base rest periods. (This is a significant difference with the airline crew scheduling case, where the credited flight time is computed as the maximum of actual flight time, a proportion of the total duty time, and another proportion of the total trip time; thus, total trip time may weigh significantly more than the total flight time.) Hence, it cannot be economically more attractive to use extended routes when normal routes may cover the same segments. Therefore, the only relevant extended routes are those made up of segments which are too long to allow the driver to return to his base on the same day. Consequently, normal routes are generated first, and extended routes are built only when required (e.g., when the regular services specified in the load plan cannot be covered by normal routes alone). These extended routes are constructed by using a column generation approach in a post-optimal stage of the solution method.

Finally, several alternatives exist to the problem of obtaining the integer optimal solution to the set covering problem, given the optimal solution to its linear relaxation. Exact algorithms (branch-and-bound, for example) may be used when exact solutions are required. When good approximate solutions are acceptable, a heuristic procedure ([20]; see also [6]) which normally produces solutions very close (less than 1%) to the optimal ones may be used. The results reported in Section 6 are obtained by using this last approach.

5.2 Generating Regular Driver Routes

To achieve our objective of a relatively small set of "good" regular driver routes, we decompose the generation procedure into several stages: normal routes, cyclic normal routes and, eventually, extended routes.

Normal driver routes are generated by a *chaining* procedure. This is an implicit enumeration scheme which chronologically scans the list of segments and combines them according to legal and efficiency criteria. Segments are suitably sorted to allow for the utilization of very efficient pruning rules (based, for example, on the maximum waiting time at a terminal, the maximum permitted working time, etc.).

Normal routes are independently generated for each day of the planning period, and for each driver base in the network. In order to allow for more flexibility in the specification of the segment departure times, some *overlapping* is permitted between the starting time of one segment and the arrival time of the previous segment in the route path. The temporal and economic characteristics of the route are, however, correctly computed; the appropriate segment starting times being adjusted if the route belongs to the optimal solution.

The chaining process stops as soon as the driver returns to his base, and the relaxed set covering subproblem is solved. Dummy routes, with high costs, are used to cover the segments uncovered otherwise. In the problems we studied, the set of normal routes was sufficiently limited to allow for all the routes to be considered by the simplex at this stage. Only cyclic routes with non negative marginal costs are generated, by combining in a post-optimal phase routes from the initial route set. Feasibility rules, marginal values and post-optimality criteria are used to limit the number of initial routes considered in the process, as well as that of new routes.

An extended driver route is composed of a sequence of segments, operations at terminals and off-base rest periods. Consequently, the same type of chaining procedure may be used to build them. Previous studies of various crew scheduling problems have shown, however, that it is much more efficient to define the path of an extended route as a sequence of duty periods and off-base rest periods. Extended routes are then generated by a two-phase procedure:

(i) Generate duty periods;

(ii) Form extended routes by legally combining duty periods and rest periods.

The main advantage of this type of approach is that only a relatively small number of "good" routes are generated. Indeed, duty periods are still not very numerous (as compared with the total number of legal routes, of course), and are normally simultaneously generated with the normal routes. The route generation phase may then be postponed until quite late in the resolution process, and may be performed by a column generation type of algorithm.

Two methods, both developed for the airline crew scheduling problem, are available for the generation of extended driver routes. Lavoie, Minoux and Odier[13] (see also [10] and references therein) generate pairings by using a shortest path algorithm on an extended space-time diagram. The nodes of the diagram correspond to the various *states* that a duty period may be in according to its position in the route path and the current cost and productivity characteristics of the route, while links indicate a legal connection between two duty periods. Crainic and Rousseau[6] take a different approach and use a chaining procedure, where strict rules (based on marginal values of pairings and duty periods) limit the access of duty periods to the generation process.

It is noteworthy that both methods require the same preliminary work (generation of duty periods), and take the same broad approach of solving a series of linear problems, each followed by computations of marginal costs and generation of routes. When using the first approach, one also has to enumerate the possible states of duty periods and build the extended diagram. This method has been shown to perform well for airline crew scheduling problems, especially when the number of states is relatively low and the number of duty periods in a route is relatively high. On the other hand, the performance of the chaining approach becomes less interesting as the length, in number of duty periods, of the route path increases.

For LTL motor carrier transportation, the number of possible states is relatively low, due to the limited number of rules concerning the driver workload (as compared with the air transport industry, for example). However, since the length of the typical route is also generally small (two days in the Canadian context), we use the chaining approach for the generation of extended regular driver routes.

6. EXPERIMENTAL RESULTS

The numerical behaviour and efficiency of the solution approach presented in the previous section has been analyzed in detail by Crainic and Roy[8]. We now describe the type of analysis that may be performed, and identify the gains that may be expected, when applying this methodology to real LTL applications.

6.1 The Data Set

Our tests were conducted on data from a large Canadian LTL motor carrier. The test network covered most of the provinces of Québec and Ontario and some parts of the Maritime provinces. The data corresponds to the set of regular driver routes that the company operated in 1988, plus information concerning the timing of actual tractor movements – day and time of departure of each route occurrence. The company works on a five-day planning period and operates a network of twenty terminals, each located in a city which is also a driver base. The available data thus reflected the company's actual choices, policies and operational logic.

Tables 2 and 3 summarize information concerning the routes operated by the company. Table 2 relates to routes. It displays the relative distribution of the 37 metaroutes and 378 driver routes that were used by the company. Table 3 presents information concerning the distribution of segments according to the number of departures during the planning period. Data are displayed for the metasegments and the corresponding segments.

Type	Metaroutes		Driver Routes	
Normal	31	84%	354	93.6%
Extended	6	16%	24	6.4%
Total	37	100%	378	100%

Table 2. Initial Route Distribution

It is worth noting that these figures justify our assertions concerning the limited usage of extended routes: more than 93% of all operated routes are made up of one duty period only. Another interesting feature is the fact that some 90% of all metasegments, which represent over 93% of all segments, have at least five departures during a planning period. In fact, over

two thirds of them (80% of the segments) have two daily departures, while one third (70% of the segments) have at least three departures daily. We also observed that, in general, the various departures of a metasegment during the planning period tend to be evenly distributed amongst the days, while departures of metasegments which "force" the use of extended routes are planned every other day. Finally, a majority of the operated bids (80% of them) prescribe that the driver take the same route at each departure. These figures amply illustrate the regularity of the operations which correspond to the stable part of demand, and empirically justify our approach.

Departures per Period	Number of Metasegments	Relative Frequency	Number of Segments	Relative Frequency
0 – 4	6	11%	24	2.71%
5 – 9	12	22%	66	7.45%
10 – 14	18	33%	184	20.77%
15 – 19	2	4%	30	3.39%
20 – 24	4	7%	88	9.93%
25 – 29	2	4%	50	5.64%
30 – 34	0	0%	0	0%
35 – 39	2	4%	72	8.13%
40 – 44	2	4%	80	9.03%
45 – 49	2	4%	90	10.16%
50 +	4	7%	202	22.80%
Total	54	100%	886	100%

Table 3. Segment Distribution in Initial Route Set

6.2 Validation

Two sets of segments were generated for validation purposes. The first one was directly extracted from the operated bids. Its segment departure times are, therefore, those indicated in company documents. For the second set of segments, we generated departure times following the procedure presented in Section 4, with special care being given to ensure that the latest departure is such that the driver still arrives on time at the destination.

Regular driver routes have been generated, and the optimal set of routes has been selected, by starting with each of these two segment sets. The route generation and selection procedures were shown to be highly efficient and to produce a limited number of routes with

"good" characteristics in terms of route composition (number and arrangement of segments), number of tractor-without-trailer (TWT) movements, number of segments covered by more than one route, etc.

Extremely interesting results may be derived by comparing the solutions obtained by our procedure with those of the initial data set, as illustrated in Table 4. Indeed, our search for efficiency translates into significant improvements in route productivity measures for roughly the same total driver wages

Criterion	1st segment set	2nd segment set
Productivity	+ 13.11%	+ 17.05%
Number of routes	- 11.60%	- 12.92%

Table 4. Comparing our results with the company's performance

The design of the routes is, in general, quite similar from one set to another, and several routes are common to all sets. This is quite natural, if one takes into account that the segments we use are rather directly derived from the routes operated by the company. Yet, one also notes that the application of our methodology produced a selection of routes which are more efficient and more productive, and which make better use of the time available during a duty period. Thus, for example, two short routes, which are generated in all sets, have been combined in our optimal solution to result in a more efficient route.

These improvements resulted in 10% fewer drivers, and therefore trucks, required to satisfy demand. Thus, the gains obtained in designing regular driver routes are directly reflected in the utilization of the equipment and in driver and rolling-stock requirements. We also believe that even more significant results may be obtained by integrating regular driver route design to the overall planning of operations, in particular to the tactical plan.

6.3 Driver Route Analysis

Our approach to generating and selecting regular driver routes is also very useful when various sensitivity or scenario analyses are required. This occurred, for example, when we were asked to evaluate the impact of recent changes made to the legislation that regulates the

permissible working hours for intercity truck drivers in the Province of Québec. Table 5 displays some of the preliminary results obtained when varying the *maximum driving time* and the *maximum duty (working) time* in a day. These results were obtained by using data from the same LTL motor carrier mentioned earlier. Note that overtime starts after eleven worked hours. Also, the average daily costs required to cover the segments are computed based on the actual costs (in Canadian Dollars), and costing formula, provided by the company. They correspond to an average of 533 driving hours to be covered daily over the planning horizon.

Maximum Duty Time per day (hours)	Maximum Driving Time per day (hours)	Average Duty Time per route (hours)	Average Driving Time per route (hours)	Average daily cost to cover the segments		
				Manpower ($/day)	Equipment ($/day)	Total ($/day)
12	8	9.60	7.32	18,140	10,254	28,394
	9	9.88	7.93	16,825	9,247	26,072
	10	9.88	8.09	15,749	8,652	24,401
15	8	9.60	7.32	18,140	10,254	28,394
	9	9.88	7.93	16,825	9,247	26,072
	10	9.92	8.17	15,462	8,458	23,920
	11	10.23	8.60	15,212	8,047	23,259
	12	10.48	8.90	15,175	7,792	22,977
	13	10.51	8.96	14,769	7,582	22,351
	14	10.59	9.06	14,801	7,530	22,331
16	13	10.51	8.96	14,769	7,582	22,35
	14	10.60	9.06	14,780	7,514	22,294
	15	10.60	9.06	14,780	7,514	22,294
17	13	10.51	8.96	14,769	7,582	22,351
	14	10.60	9.06	14,780	7,514	22,294
	15	10.61	9.08	14,790	7,504	22,294
	16	11.99	10.40	15,281	6,550	21,831

Table 5. Scenario Analysis on Driving Time Regulations

First, we observe that route productivity, as measured by the average duty and driving times per route, increases steadily with the maximum number of driving hours allowed on a daily basis. This comes as no surprise, since our approach clearly favours the selection of productive routes, i.e. those with a relatively high number of driving hours, within the permissible company and government rules and regulations. This could also explain why our results are much more sensitive to variations in the maximum driving time than they are to changes in the maximum duty time.

We have also found that the various cost results presented in Table 5 are consistent with what could be expected. Indeed, the average total daily cost required to cover the segments (which represent the stable part of demand for a given period) decreases steadily as driving and duty time restrictions are relaxed.

It is also interesting to note that the new regulation on driving hours in the Province of Québec restricts the maximum duty and driving times to 15 and 13 hours per day respectively. Referring to the results displayed in Table 5, we can conclude that this new regulation will not have a major economic impact on the motor carrier that provided us with data. This would not be the case, however, if the Québec government was to adopt the same rules as those applied in the United States where maximum driving time is restricted to 10 hours. Our analysis indicates that such a change would result in a daily cost increase of some 1569 $CND, (if the maximum duty time remained at 15 hours per day), or approximately 400,000 $CND per year for this particular company.

CONCLUSION

We have presented an integrated approach for the problem of designing regular driver routes in the less-than-truckload motor carrier industry. It represents one of the first comprehensive research efforts in this area, and its success is both rewarding and encouraging. Further research is however needed to explore alternate modeling approaches that could deal more directly with the fluidity of the temporal characteristics (e.g., truck departures obey schedules which are significantly less fixed than airline ones) of the problem.

Yet, results obtained thus far clearly indicate that the potential savings which may be obtained by using the approach presented herein are quite substantial. Our methodology can provide assistance to motor carriers in making better use of their human and material resources, and hence enhance the overall planning process within the company and improve their productivity and competitivity.

Finally, we plan to integrate the models and algorithms presented in this paper into a comprehensive interactive graphics decision-support system. In particular, we have allowed for the possibility of user interaction at each stage in the driver route design process (see Figure 2). This would enable users to manage the overall process and contribute their special experience and knowledge of the problem; they could, for example, bring various unquantifiable criteria into the process, or impose certain personal preferences which otherwise would go unobserved. Past experience has also shown that such a presentation of operations research methodology facilitates its penetration into industrial practice and enhances its chances of wide utilization.

ACKNOWLEDGMENTS

We wish to thank Ms. Lucie Bibeau for her excellent assistance in programming, and Mr. Alain Blanchette for his consistent and successful efforts in providing us with a very effective communication link with the industry. We also wish to thank the Editor, the Associate Editor and the referees for their thorough review of the first version of the manuscript; their suggestions and comments have helped us write a better paper. Finally, we express our gratitude to the following funding organizations: Le Fonds FCAR Québec, C.R.S.N.G. Canada, le Ministère des transports du Québec, and the Foundation of the Université du Québec à Montréal.

BIBLIOGRAPHY

1. C. Barnhart, E. Johnson and R. Anbil, "A Column Generation technique for the Long-Haul Crew Assignment Problem", Research Report, School of Industrial and Systems Engineering, Georgia Tech, 1991.

2. A. Blanchette, "Etablissement des routes et horaires réguliers dans les entreprises de transport routier des marchandises", M.B.A. Thesis, département des Sciences administratives, Université du Québec à Montréal, Publication #663, Centre de recherche sur les transports, Université de Montréal, 1989.

3. T.G. Crainic, "Un modèle de planification tactique pour le transport ferroviaire des marchandises", Ph.D. Thesis, Publication #247, Centre de recherche sur les transports, Université de Montréal, 1982.

4. T.G. Crainic, "Rail Tactical Planning: Issues, Models and Tools", in *Freight Transport Planning and Logistics*, L. Bianco and A. La Bella (Eds.), 463-509, Springer-Verlag, Amsterdam, 1988.

5. T.G. Crainic and J.-M. Rousseau, "Multicommodity, Multimode Freight Transportation: A General Modeling and Algorithmic Framework for the Service Network Design Problem", *Trans. Res.* **20**, 225-242 (1986).

6. T.G. Crainic and J.-M. Rousseau, "The Column Generation Principle and the Airline Crew Scheduling Problem", *INFOR* **25**, 136-151 (1987).

7. T.G. Crainic and J. Roy, "O.R. Tools for Tactical Freight Transportation Planning", *EJOR* **33**, 290-297 (1988).

8. T.G. Crainic and J. Roy, "Une approche de recouvrement d'ensembles pour l'établissement d'horaires de chauffeurs dans le transport routier de charges partielles", *R.A.I.R.O. – Recherche Opérationnelle* **2**, 1-36 (1990).

9. M. Desrochers and F. Soumis, "A Column Generation Approach to the Urban Transit Crew Scheduling Problem", *Trans. Sci.* **23**, 1-13 (1989).

10. J. Desrosiers, Y. Dumas, M. Desrochers, F. Soumis, B. Sanso and P. Trudeau, "A Breakthrough in Airline Crew Scheduling", Report G-91-11, GERAD, Montréal, Canada, 1991.

11. M.L. Fisher and K. Pradeep, "Optimal Solution of Set Covering/Partitioning Problems Using Dual Heuristics", *Man. Sci.* **36**, 674-688 (1990).

12. I. Gershkoff, "Optimizing Flight Crew Schedules", *Interfaces* **19**, 29-43 (1989).

13. S. Lavoie, M. Minoux and E. Odier, "A New Approach for Crew Pairing Problems by Column Generation with an Application to Air Transportation", *EJOR* **35**, 45-58 (1988).

14. W.B. Powell, "A Local Improvement Heuristic for the Design of Less-Than-Truckload Motor Carrier Networks", *Trans. Sci.* **20**, 246-257 (1986).

15. W.B. Powell and Y. Sheffi, "The Load Planning Problem of Motor Carriers: Problem Description and Proposed Solution Approach", *Trans. Res.* **17**, 471-480 (1983).

16. W.B. Powell and Y. Sheffi, "Design and Implementation of an Interactive Optimization System for the Network Design in the Motor Carrier Industry", *Opns. Res.* **37**, 12-29 (1989).

17. J.-M. Rousseau, "Crew Scheduling Methods in the Operation and Planning of Transportation Systems", in *Transportation Planning Models*, M. Florian (Ed.), 439-474, North-Holland, 1984.

18. J. Roy, "Un modèle de planification globale pour le transport routier des marchandises" , Ph.D. Thesis, Publication #402, Centre de recherche sur les transports, Université de Montréal, 1984.

19. J. Roy and L. Delorme, "NETPLAN: A Network Optimization Model for Tactical Planning in the Less-than-Truckload Motor-Carrier Industry", *INFOR* **27**, 22-35 (1989).

20. H.M. Salkin and K. Mathur K., "Fundamentals of Integer Programming", North-Holland, New York, U.S.A., 1989.

APPENDIX. COST MEASURES FOR DRIVER ROUTES

We examined several possible formulae for computing route costs which achieve the objectives discussed in Section 3. In particular, we considered the following hypotheses (defined here for normal routes only):

H1. The route cost equals its *average hourly cost*: $c_r^1 = \gamma_r / w_r$.

H2. Compute the route's *average unit cost* $c_r^m = \{K^m + c_r^n (\omega^n - \omega^m) + c_r^o (\omega^o - \omega^n)\}/\omega^o$, over the maximum duty time length $[0, \omega^o]$, and define the route cost as $c_r^2 = c_r^m w_r$.

H3. Let K^i be an hourly rate for the trip *idle time*, defined as the difference between the driving time and the normal amplitude of a duty period, and K^o a "penalty" for working overtime. The route cost, c_r^3, is then $\gamma_r + K^i (\omega^n - w_r)$, if $w_r \leq \omega^n$, and $\gamma_r + K^o (w_r - \omega^n)$, otherwise.

To briefly illustrate the behaviour of these formulae, we give in Table 6 the results obtained for six hypothetical routes. The total trip and driving times are indicated for each route. The following parameters are used: $K^m = 90$, $\omega^m = 9$, $\omega^n = 12$, $\omega^o = 14$, $c_r^n = 10$, $c_r^o = 15$ (although only indicative, these values are representative of actual practice). For the third cost formula, we first (column *(a)*) use $K^i = 5$ and $K^0 = 0$, and then (column *(b)*) $K^i = K^o = 15$. The c_r values are computed by using the formula defined in Section 3.

Route	w_r	δ_r	γ_r	c_r^1	c_r^2	c_r^3(a)	c_r^3(b)	c_r
1	5hrs	4hrs	90.00	18.00	53.57	125.00	195.00	5.00
2	7hrs	4hrs	90.00	12.86	75.00	115.00	165.00	2.00
3	10hrs	7hrs	100.00	10.00	107.14	110.00	130.00	3.00
4	11hrs	7hrs	110.00	10.00	117.85	115.00	125.00	4.00
5	13hrs	10hrs	135.00	13.50	139.28	135.00	345.00	3.50
6	14hrs	10hrs	150.00	15.00	150.00	150.00	570.00	5.00

Table 6. Numerical Behaviour of Various Route Cost Measures

The first two hypotheses, and their variants, are quite simple, even simplistic, and require only minimal computing efforts. The route discrimination issue is only partially addressed, while the productivity issue is not addressed at all.

For the third formula, one assumes that the normal duty period amplitude is a worthy goal for a productive driver route. This type of cost formula yields values which are not far from the actual wages paid to drivers and it also permits to penalize short and, when so desired, overtime routes. Moreover, the *idle time cost*, K^i, offers interesting opportunities for analysis concerning the relative value of the various activities that make up driver routes. The calibration of the idle time cost represents, however, a very serious problem, since its value has a significant impact of the route costs and, therefore, on the final route choice, as illustrated in Table 6. One may note the insidious effect of the idle time cost value on the relative cost ordering for routes with duty times between the minimum guaranteed and the normal spread. Besides, this type of formulae does not adequately address the productivity issue, since two routes of equal total trip time will still receive the same cost, irrespective of their actual driving time.

Hence, we have chosen the formula defined in Section 3, which adequately reflects the composition and productivity of regular driver routes.